arXiv:1404.4540v1 [cs.SI] 17 Apr 2014

Collective computation in a network with distributed information

A. Córdoba
Departamento de Física de la Materia Condensada, Universidad de Sevilla,
P. O. Box 1065, 41080 Sevilla, Spain
cordoba@us.es
D. Aguilar-Hidalgo
Max Planck Institute for the Physics of Complex Systems,
Nöthnitzer Str. 38, 01187 - Dresden, Germany
M. C. Lemos
Departamento de Física de la Materia Condensada, Universidad de Sevilla,
P. O. Box 1065, 41080 Sevilla, Spain

Abstract

We analyze a distributed information network in which each node has access to the information contained in a limited set of nodes (its neighborhood) at a given time. A collective computation is carried out in which each node calculates a value that implies all information contained in the network (in our case, the average value of a variable that can take different values in each network node). The neighborhoods can change dynamically by exchanging neighbors with other nodes. The results of this collective calculation show rapid convergence and good scalability with the network size. These results are compared with those of a fixed network arranged as a square lattice, in which the number of rounds to achieve a given accuracy is very high when the size of the network increases. The results for the evolving networks are interpreted in light of the properties of complex networks and are directly relevant to the diameter and characteristic path length of the networks, which seem to express "small world" properties.

1 Introduction

We propose a model for a distributed information network in which each node at a given time can obtain information from a limited number of other nodes to which it is connected. From the information residing in these nodes, each one of them can perform specific tasks (calculation, classification, etc.) involving all the information contained in the network nodes. In a system where information is distributed at different sites, the access to this information can be an expensive procedure if the number of sites is high. Moreover, if a set of N nodes want to do at any given time a calculation of a magnitude involving information contained in all other ones, the direct access of all to all nodes requires a number of requests of order N^2, which is very high if N is large.

Here we pose the problem of a set of N elements (nodes), each characterized by the value of a magnitude s, in general different for each node, so that each and every one of the elements wish to calculate the average value of s (or a function of s) in the set, by accessing at one time to information contained in other q elements of the set ($q << N$) with which it is connected. This is in line with systems using epidemic protocols [1, 2, 3] or gossip ones [4, 5, 6], such as the newscast protocol [7]. In these systems the goal is not to enable both point-to-point communication between nodes, but rapid and efficient dissemination of information. The system intends to perform a specific collective task (for example, calculating the average value of a variable in the set of nodes, setting the position of each node in a ranking according to the set value of the variable, etc.) so that, eventually, all nodes have access to the result obtained from the set. To do this we consider a network in which each element is connected to other q ones (neighbors), from which it extracts information. The neighborhood of each element can be changed along the process of calculation, so that each node can exchange a neighbor randomly with another node. We discuss issues such as scalability and convergence, considering different sizes and data sets. We also compare the results obtained for the dynamically changing network with those obtained from the fixed static network. To do this we compare network configurations from different times of evolution with those corresponding to the initial fixed network (which may take the form of a conventional cellular automaton).

2 The model

In a gossip framework there is an exchange of information among system elements In this exchange of information, which is a dynamic process, one receives information from another (also can be a reciprocal exchange). In turn, the receiver of information can give information to other peers. Overall the transmission process basically comprises three aspects: peer selection, data exchanged or transmitted between peers and data processing. In our model we consider a network arranged in a two-dimensional lattice (we do this initially to compare with a conventional cellular automaton with Moore's neighborhood [8]) with $m \times n$ sites forming a square lattice. To each node is randomly assigned the value of a variable s, in general different for each node. The objective is that each node "knows" the average value of this variable in the whole at the end of the process. Each node has access to limited information on each step of calculation. To obtain specific results, it has been considered that each node, at each instant, can only store eight values of other nodes, and that initially each node has access to the data of the eight nearest neighbors in the lattice. This can be matched with a two-dimensional cellular automaton with Moore's neighborhood. Throughout the process the neighborhood of each site is not fixed, but each node exchanges a neighbor with another randomly chosen node. This implies a double selection (with appropriate calls to random numbers routines), one of them for the site of exchanging and the other for choosing the neighbor exchanged. The average value of the variable of each site and its eight neighbors (inputs) is assigned to the variable of the corresponding site for each system update. The updating of all sites is simultaneous, although it could also be carried out sequentially. Since in the course of evolution neighborhoods randomly change, the initial regular structure of the network is irrelevant (only holds the fact that each node has eight input and eight outputs, generally different). However, in order to formally create a by analogy with the above mentioned cellular automaton, we initially identify each node as a point on a square lattice $m \times n$ (Figure 1A), but now, after the dynamic exchange of information, the system is turned into a complex directed network and connections do not generate a square

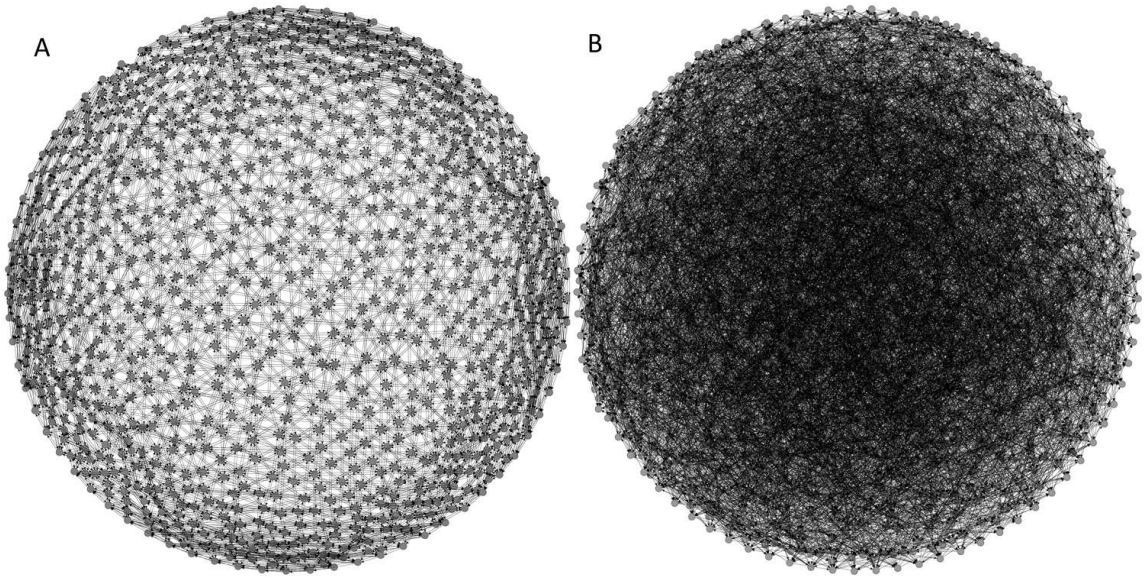

Figure 1: Representation of 32x32 networks (1024 nodes and 8192 links). A) Network generated by the Cellular Automaton with Moore's neighborhood and used as an initial condition for the Evolutionary method. B) Evolutionary network representation. This last network is far less ordered than the cellular automaton case.

lattice (Figure 1B)[1].

To test the scalability we have considered various system sizes (32×32, 100×100, 320×320 and 1000×1000). Also we have considered two data distributions: (a) a uniform random distribution, (b) a distribution of values grouped into four quadrants as shown schematically in Figure 4. For comparison, as mentioned above, three cases are considered: the network that dynamically changes, the fixed network with the configuration that is reached at a given moment of the dynamical evolution and the fixed network with the initial square lattice structure (cellular automaton). In the latter two cases, the neighborhoods do not change and each node has permanent access to its pre-established eight neighbors.

We defined a parameter, b, to evaluate how the computation progresses in each system update, i.e., convergence to the desired value with a given accuracy.

$$b = \text{standard deviation/mean value} \tag{1}$$

We have imposed the condition that the calculation is stopped when the variation of b between two successive updates is less than a preset value. We have also established a limited number of system updates. Moreover, using the Cytoscape software [9, 10], we have analyzed some of the properties of the network [11-12] to establish the relationship between its topology and the degree of computation efficiency.

[1]Figures 1 A and B are represented in a circular layout in order to compact the graph, as an orthogonal layout (square lattice like) is not friendly looking due to the high number of nodes.

Evolutionary Network					
$m \times n$	$b < 10^{-1}$	$b < 10^{-2}$	$b < 10^{-3}$	$b < 10^{-4}$	$b < 10^{-5}$
$32 \times 32 \sim 10^3$	3	6	8	11	13
$100 \times 100 \sim 10^4$	3	6	9	11	14
$320 \times 320 \sim 10^5$	3	6	9	11	14
$1000 \times 1000 \sim 10^6$	3	6	9	11	14

Table 1: Number of updates required to reach the value of b in a evolutionary network for a set of randomly distributed data.

Fixed networks Configuration of evolutionary network when it reaches: (i) $b < 10^{-1}$, (ii) $b < 10^{-5}$.					
$m \times n$	$b < 10^{-1}$	$b < 10^{-2}$	$b < 10^{-3}$	$b < 10^{-4}$	$b < 10^{-5}$
$32 \times 32(i)$	2	5	8	12	16
$320 \times 320(i)$	2	5	9	13	17
$32 \times 32(ii)$	2	4	7	9	12
$320 \times 320(ii)$	2	5	7	10	12

Table 2: Number of updates required to reach the value of b in two fixed network for a set of randomly distributed data.

3 Results and Discussion

We have taken the value of $b = 10^{-5}$ as the its limit of b and have scored the number of updates required to reach a value of b less than the successive powers of 10 from 10^{-1} to 10^{-5}. The Table 1 shows the results for the evolving network for a set of data randomly distributed according to a uniform distribution. As can be seen, the number of updates is low and the scalability is very good, because for large variations in the size of the system the number of updates nearly remains the same. Each node has to perform a small number of updates. Since for each update of the entire system one must perform $N = m \times n$ updates of the values of the nodes, then, the total number of individual operations is virtually proportional to system size, i.e. escalates as N (a direct calculation escalates as N^2). Figure 2A graphically show this behavior where the number of updates slowly increases when b is diminished. It must be noticed that the number of updates remains the same for a certain b independently of the network size.

Let's now change the computation order. Instead of updating the system as the network evolves, the evolution takes first place and then, for a fixed network successive updates are done. The results for this case can be seen in Table 2. Here the configurations that are considered are those reached with the evolutionary network when (i) $b < 10^{-1}$ and (ii) $b < 10^{-5}$. In these cases, the differences with the evolutionary network are not very significant, but for larger values of the limit of b convergence is slightly faster, and for lower values of this limit convergence is slightly slower in the case (i) and slightly faster in the case (ii). Again, the number of updates slowly increases linearly when b is diminished in both cases (i) and (ii) (Figure 2B).

Table 3 shows the results obtained for the same random distribution of data using the fixed network forming a square lattice (cellular automaton with Moore's neighborhood). As can be seen, the number of updates required to achieve a given accuracy rapidly grows with size. To be concrete, the number of updates increases exponentially when b is diminishes (Figure 2C) and the size of the

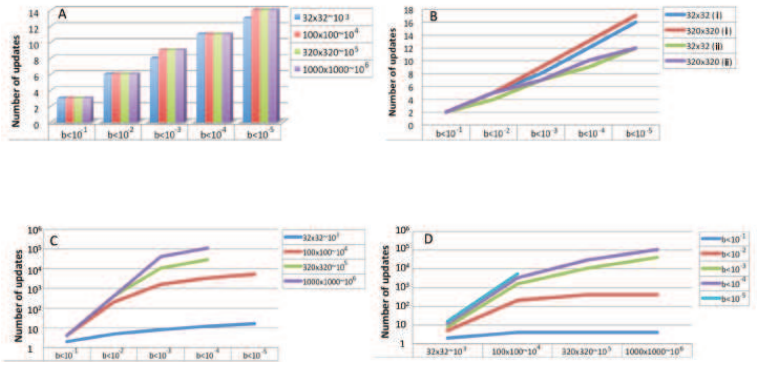

Figure 2: (A) Number of updates required to reach the b value in an evolutionary network for a randomly distributed data set. The bar graph shows a linear increase of the number of updates with the required precision and remains constant when increasing the network size. (B) Number of updates required to reach the b value in two fixed networks for a randomly distributed data set: (i) Evolutionary network configuration when $b < 10^{-1}$. (ii) Evolutionary network configuration when $b < 10^{-5}$. As in the case of using evolutionary network, the number of updates needed to reach a certain precision in the calculus increases linearly with this precision. In this case, the increase in the network size makes the number of updates to remain nearly constant. (C, D) Number of updates required to reach the b value applying a cellular automaton with Moore neighborhood for a randomly distributed data set. (C) varying b. (D) Varying the network size. Contrary to the two anterior cases of study, the use of cellular automata makes the number updates increase in an exponential tendency as the precision and the network size grow.

Cellular automaton on square lattice with Moore's neighborhood					
$m \times n$	$b < 10^{-1}$	$b < 10^{-2}$	$b < 10^{-3}$	$b < 10^{-4}$	$b < 10^{-5}$
$32 \times 32 \sim 10^3$	2	5	8	12	16
$100 \times 100 \sim 10^4$	4	206	1608	3350	5255
$320 \times 320 \sim 10^5$	4	419	10678	29568	(*)
$1000 \times 1000 \sim 10^6$	4	420	40255	109068	(*)

Table 3: Number of updates required to reach the value of b by applying a cellular automaton with Moore's neighborhood for a set of randomized data. (*) The stop criterion is reached before obtaining the value of b indicated.

network increases (Figure 2D).

To analyze the relationship of these results with the network structure, we have made an analysis of some of the properties of these networks. We have considered the clustering coefficient, the diameter and the characteristic path length (CPL). The clustering coefficient of a node is the ratio p/r, where p is the number of links between the neighboring nodes and r is the maximum number of links that would be possible among them; the clustering coefficient of the network is the average value of clustering coefficients of all the network nodes. The network diameter is the maximum distance between two nodes. The characteristic path length gives the expected distance between any two nodes. It is defined as the average number of steps along the shortest paths for all possible pairs of network nodes. As can be seen, the shorter the CPL the better the communication along the network. By construction the number of links of each node in these networks is the same (eight inputs and eight outputs). The clustering coefficient of the evolutionary network has a very low value and that of the cellular automaton is high. The two significant parameters are the diameter of the network and the characteristic path length. The values of these parameters for the fixed networks considered are shown in Table 4. As can be seen, whereas in the four evolutionary network configurations considered the diameter and the characteristic path length have small values, and do not change significantly with the limit of b, in the cellular automaton these values are much higher and rapidly growing with size (Figure 3A). For the cellular automaton convergence is very slow, the number of updates required to reach the limits of b is very high and the scalability is very poor. Therefore this fact suggests that the speed of convergence and scalability of the calculus in the network is strongly associated with the values of these two parameters. This can be seen as manifesting the property of "small world", which appears in different types of complex networks [11, 12]. Topologically speaking, the evolutionary networks are distributed in a more sparse way than the cellular automaton, according to the clustering coefficient values. Regarding CPL and diameter values, the evolutionary networks are much 'better' connected than the cellular automaton, in terms of information dissemination. A network with low CPL means that the information contained in one node is more accessible to the rest of the nodes than in a network with high CPL. All this gives a method that re-distributes a network in a way where information is easy to share, and so, computationally speaking, with a low cost in collective computations.

Next we have examined the influence of the way in which data are distributed. Instead of a random distribution, we have considered other one in which data is strongly grouped. This consists of four different values distributed in the nodes associated to each of the four quadrants shown in Figure 4. The results for the evolutionary network are shown in Table 5. As can be seen scalability is very good as in the case of the data with random distribution, although the number of updates

Topological Parameters.			
Network	Diameter	Characteristic path length	Clustering coefficient
Evolutionary network 32×32 when reaching $b < 10^{-3}$	5	3.5	0.007
Evolutionary network 100×100 when reaching $b < 10^{-3}$	6	4.6	0.001
Evolutionary network 320×320 when reaching $b < 10^{-1}$	8	5.9	0.036
Evolutionary network 320×320 when reaching $b < 10^{-5}$	8	5.7	0.000
Cellular automaton 32×32	16	10.7	0.429
Cellular automaton 100×100	50	33.3	0.429
Cellular automaton 320×320	160	106.7	0.429

Table 4: Values of some characteristic parameters of the listed networks.

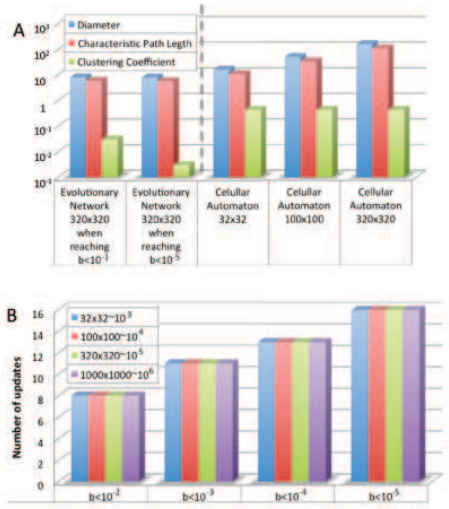

Figure 3: (A) Topological parameters measured in the analyzed networks. The bar graph shows a constant behavior for the diameter and the characteristic path length in the evolutionary networks (at the left side of the vertical discontinuous line). In the cellular automata case (right side of the vertical discontinuous line), these parameters highly increase with the network size. The clustering coefficient has a low value in the evolutionary network and higher in the cellular automata. Though the connectivity remains always stationary, the evolutionary networks present a much more sparse topology than in the cellular automata cases, which is indicative for a more efficient connectivity in terms of information dispersion. (B) Number of updates required to reach the b value in a specific case using evolutionary networks when the data is not randomly distributed but strongly grouped. In this case the scalability is still very good as in the case with randomly distributed data. The number of updates also slowly increases in a linear way as in the rest of studied cases for the evolutionary network.

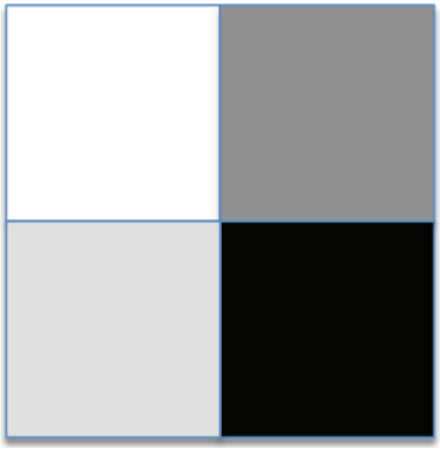

Figure 4: Schema of data grouped by quadrants.

Evolutionary network				
$m \times n$	$b < 10^{-2}$	$b < 10^{-3}$	$b < 10^{-4}$	$b < 10^{-5}$
$32 \times 32 \sim 10^3$	8	11	13	16
$100 \times 100 \sim 10^4$	8	11	13	16
$320 \times 320 \sim 10^5$	8	11	13	16
$1000 \times 1000 \sim 10^6$	8	11	13	16

Table 5: Number of updates for an evolutionary network with a set of tightly grouped data.

required is somewhat greater than in the case of random data distribution. As in the rest of the studied cases, the number of updates slowly increases in a linear way (Figure 3B).

Finally, note that if during the process you add new nodes are added or deleted, the update could proceed as follows:

- If a node is deleted, its input links could be redirected to its output target nodes, so that those nodes that had access to it now are addressed to those to which the deleted node accessed.
- If a node is inserted, it provides access to eight of the existing nodes (as input links), and in exchange, each of these nodes will deliver to the new node a link to one of its neighbors (as output nodes).

This will allow to perform a system robustness analysis. If the variation in the number of nodes is not very sudden, it is expected the overall system to be not severely impacted, given the information transmission speed. Even in the case of a change in many nodes, it can be expected that a good accuracy in the result of the global parameter will be reached again in a reduced number of updates, i.e. resilience is very good.

4 Conclusions

In a networked system, which nodes have a limited access to information only from a few neighbors, the collective computation involving the whole data set is very efficient when the network changes dynamically, or with a fixed structure generated by the same dynamic process. In this case, a fast convergence towards the average value is achieved. Also the number of computing updates shows a good scalability with the size of the network. This is associated with the diameter and the characteristic path length of the network. Given the high rate with which the system converges to the desired value, the system is also is expected to be robust to changes in the number of nodes or in the values distribution assigned to the network. These results clearly contrast with what happens in a fixed regular network in which the convergence is slow and the computation required hugely grows when the size of the system increases.

Acknowledgement

This work is partially financed by the Project FIS2008-04120 of the Spanish Ministry for Science and Innovation (MICINN).

References

[1] P. T. Eugster, R. Guerrraoui, A.-M, Kermarrec, L. Massouli, Epidemic information dissemination in distributed systems, IEEE Computer 37 (2004) 60–67.

[2] P. De, S. K. Das, Epidemic models, algorithms and protocols in wireless sensor and ad-hoc networks, in *Algorithms and protocols for wireless sensor networks*, John Wiley and Sons, 2008.

[3] C. Anagnostopoulos, S. Hadjiefthymiades, E. Zervas, An analytical model for multi-epidemic information dissemination, J. Parallel Distrib. Comput. 71 (2011) 87–104.

[4] A. Kermarrec, M. V. Steen, Gossiping in distributed systems, Open. Syst. Rev. 41 (2007) 2–7.

[5] M. H. Islam, S. Waheed, I. Zubai, An efficient gossip based overlay network for peer-to-peer Networks, First International Conference on Ubiquitous and Future Networks (2009) 62–67.

[6] S. Tang, E. Jaho, I. Stavrakakis, I. Koukoutsidis, P. V. Mieghen, Modeling gossip-based content dissemination and search in distributed networking, Computer Communications 34 (2011) 765–779.

[7] S. Voulgaris, M. Jelasity, M. van Steen, A robust scalable peer-to-peer gossiping protocol, in Agent and Peer-to-Peer Computing, G. Moro, C. Sartori, M. P. Sing, eds., Springer, 2004.

[8] S. Wolfram, A new kind of science, Wolfram Media, 2002.

[9] P. Shannon, A. Markiel, O. Ozier, N. S. Baliga, J. T. Wang, D. Ramage, N. Amin, B. Schwikowski, T. Ideker, Cytoscape: A Software Environment for Integrated Models of Biomolecular Interaction Networks, Genome Res. 13 (2003) 2498–2504.

[10] Y. Assenov, F. Ramírez, S. Schelhorn, T. Lengauer, M. Albrecht, Computing topological parameters of biological networks, Bioinformatics 24 (2) (2008) 282–284.

[11] R. Albert, A.-L. Barabasi, Statistical mechanics of complex networks, Rev. Mod. Phys. 74 (2002) 47–97.

[12] S. Boccaletti, V. Latora, Y. Moreno, M. Chavez, D.-U. Hwang, Complex networks: structure and dynamics, Physics Reports 424 (2006) 175–308.